

Narrow progressions in the primes

Terence Tao and Tamar Ziegler

Abstract In a previous paper of the authors, we showed that for any polynomials $P_1, \dots, P_k \in \mathbb{Z}[\mathbf{m}]$ with $P_1(0) = \dots = P_k(0)$ and any subset A of the primes in $[N] = \{1, \dots, N\}$ of relative density at least $\delta > 0$, one can find a “polynomial progression” $a + P_1(r), \dots, a + P_k(r)$ in A with $0 < |r| \leq N^{o(1)}$, if N is sufficiently large depending on k, P_1, \dots, P_k and δ . In this paper we shorten the size of this progression to $0 < |r| \leq \log^L N$, where L depends on k, P_1, \dots, P_k and δ . In the linear case $P_i = (i - 1)\mathbf{m}$, we can take L independent of δ . The main new ingredient is the use of the densification method of Conlon, Fox, and Zhao to avoid having to directly correlate the enveloping sieve with dual functions of unbounded functions.

1 Introduction

1.1 Previous results

We begin by recalling the well-known theorem of Szemerédi [18] on arithmetic progressions, which we phrase as follows:

Theorem 1 (Szemerédi’s theorem). *Let $k \geq 1$ and $\delta > 0$, and suppose that N is sufficiently large depending on k, δ . Then any subset A of $[N] := \{n \in \mathbb{Z} : 1 \leq n \leq N\}$ with cardinality $|A| \geq \delta N$ will contain at least one arithmetic progression $a, a + r, \dots, a + (k - 1)r$ of length k , with $r > 0$.*

Terence Tao
UCLA Department of Mathematics, 405 Hilgard Ave, Los Angeles CA 90095, USA, e-mail: tao@math.ucla.edu

Tamar Ziegler
Einstein Institute of Mathematics, Edmond J. Safra Campus, Givat Ram, The Hebrew University of Jerusalem, Jerusalem, 91904, Israel e-mail: tamarz@math.huji.ac.il

In fact, by partitioning $[N]$ into intervals of some sufficiently large but constant size $L = L(k, \delta)$ and using the pigeonhole principle, one can ensure that the progression described above is “narrow” in the sense that $r \leq L(k, \delta)$.

In [10], Szemerédi’s theorem was relativized to the primes $\mathbb{P} = \{2, 3, 5, 7, \dots\}$:

Theorem 2 (Szemerédi’s theorem in the primes). *Let $k \geq 2$ and $\delta > 0$, and suppose that N is sufficiently large depending on k, δ . Then any subset A of $[N] \cap \mathbb{P}$ with $|A| \geq \delta |[N] \cap \mathbb{P}|$ will contain at least one arithmetic progression $a, a+r, \dots, a+(k-1)r$ of length k , with $r > 0$.*

In particular, the primes contain arbitrarily long arithmetic progressions.

The proof of Theorem 2 does not place a bound on r beyond the trivial bound $r \leq N$. In contrast with Theorem 1, one cannot hope here to make the step size r of the progression as short as $L(k, \delta)$. Indeed, as observed in [1], if N is large enough then from [8, Theorem 3] one can find a subset A of $[N] \cap \mathbb{P}$ with $|A| \geq \frac{1}{2} |[N] \cap \mathbb{P}|$ (say) such that the gap between any two consecutive elements of A is $\geq c \log N$ for some absolute constant $c > 0$, which of course implies in this case that r must be at least $c \log N$ as well. Indeed, one can improve this lower bound to $\log^{k-1} N$ by a small modification of the argument:

Proposition 1. *Let $k \geq 2$, let $\varepsilon > 0$ be sufficiently small depending on k , and suppose that N is sufficiently large depending on k . Then there exists a subset A of $[N] \cap \mathbb{P}$ with $|A| \gg |[N] \cap \mathbb{P}|$ such that A does not contain any arithmetic progression $a, a+r, \dots, a+(k-1)r$ of length k with $0 < r < \varepsilon \log^{k-1} N$.*

Proof. For any $0 < r < \varepsilon \log^{k-1} N$, let N_r denote the number of $a \in [N]$ such that $a, a+r, \dots, a+(k-1)r$ are all prime. By the union bound and the prime number theorem, it will suffice to show that

$$\sum_{0 < r < \varepsilon \log^{k-1} N} N_r \leq \frac{1}{2} \frac{N}{\log N}$$

(say), since one can then form A by removing all the elements a in $[N] \cap \mathbb{P}$ that are associated to one of the N_r . But from standard sieve theoretic bounds (see e.g. [12, Theorem 5.7]) we have

$$N_r \leq C_k \mathfrak{G}(k, r) \frac{N}{\log^k N}$$

where C_k depends only on k and $\mathfrak{G}(k, r)$ is the singular series¹

$$\mathfrak{G}(k, r) = \prod_p \left(1 - \frac{1}{p}\right)^k \left(1 - \frac{v_p(r)}{p}\right) \quad (1)$$

and $v_p(r)$ is the number of residue classes $a \in \mathbb{Z}/p\mathbb{Z}$ such that at least one of $a, a+r, \dots, a+(k-1)r$ is equal to 0 mod p , so it suffices to show that

¹ All sums and products over p in this paper are understood to be ranging over primes.

$$\sum_{0 < r < M} \mathfrak{G}(k, r) \leq C_k M$$

for any $M \geq 1$ (we allow C_k to represent a different k -dependent constant from line to line, or even within the same line). One could obtain precise asymptotics on the left-hand side using the calculations of Gallagher [7], but we can obtain a crude upper bound that suffices as follows. From (1) we see that

$$\mathfrak{G}(k, r) \leq C_k \exp \left(C_k \sum_{p|r} \frac{1}{p} \right)$$

and hence by [19, Lemma E.1]

$$\mathfrak{G}(k, r) \leq C_k \sum_{p|r} \frac{\log^{C_k} p}{p}$$

and thus

$$\sum_{0 < r < M} \mathfrak{G}(k, r) \leq C_k \sum_p \frac{\log^{C_k} p}{p} \frac{M}{p} \leq C_k M$$

as required.

In the converse direction, if we use the ‘‘Cramér random model’’ of approximating the primes $\mathbb{P} \cap [N]$ by a random subset of $[N]$ of density $1/\log N$, we can asymptotically almost surely match this lower bound, thanks to the work of Conlon-Gowers [6] and Schacht [17]:

Proposition 2. *Let $k \geq 2$ and $\delta, \varepsilon > 0$, let $C > 0$ be sufficiently large depending on δ, k , and suppose that N is sufficiently large depending on k, δ, ε . Let $P \subset [N]$ be chosen randomly, such that each $n \in [N]$ lies in P with an independent probability of $1/\log N$. Then with probability at least $1 - \varepsilon$, every subset A of P with $|A| \geq \delta|P|$ will contain an arithmetic progression $a, a + r, \dots, a + (k - 1)r$ of length k with $0 < r \leq C \log^{k-1} N$.*

We remark that a modification of the argument in Proposition 1 shows that we cannot replace the large constant C here by an arbitrarily small constant $c > 0$.

Proof. We partition $[N]$ into intervals I_1, \dots, I_m of length between $\frac{C}{2} \log^{k-1} N$ and $C \log^{k-1} N$, thus $m \leq \frac{2N}{C \log^{k-1} N}$. For each interval I_i , we see from [17, Theorem 2.2] or [6, Theorem 1.12] that with probability at least $1 - \frac{\delta \varepsilon}{10}$, every subset A_i of $P \cap I_i$ with $|A_i| \geq \frac{\delta}{2} |P \cap I_i|$ will contain an arithmetic progression of length at least k . Call an interval I_i *bad* if this property does not hold, thus each I_i is bad with probability at most $\delta \varepsilon / 10$. By linearity of expectation followed by Markov’s inequality, we conclude that with probability at least $1 - \varepsilon$, at most $\frac{\delta N}{5C \log^{k-1} N}$ of the I_i are bad. Then if $A \subset P$ is such that $|A| \geq \delta|P|$, then at most $\frac{\delta}{2}|P|$ of the elements of A are contained in bad intervals, so from the pigeonhole principle there exists a good interval I_i such that $|A \cap I_i| \geq \frac{\delta}{2} |P \cap I_i|$, and the claim follows.

It is thus natural to conjecture that in Theorem 2 one can take r to be as small as $\log^{k-1+o(1)} N$.

Remark 1. If one seeks progressions inside the full set \mathbb{P} of primes, rather than of dense subsets of the primes, then the Hardy-Littlewood prime tuples conjecture [13] predicts that one can take r to be of size $O_k(1)$; indeed, one should be able to take r to be the product of all the primes less than or equal to k . In the case $k = 2$, the claim that one can take $r = O(1)$ amounts to finding infinitely many bounded gaps between primes, a claim that was only recently established by Zhang [20]. For higher k , the claim $r = O_k(1)$ appears to currently be out of reach of known methods; the best known result in this direction, due to Maynard [15] (and also independently in unpublished work of the first author), shows that for any sufficiently large $R > 1$, there exist infinitely many intervals of natural numbers of length R that contain $\geq c \log R$ primes for some absolute constant $c > 0$, but this is too sparse a set of primes to expect to find length k progressions for any $k \geq 3$.

We now consider generalizations of the above results, in which arithmetic progressions $a, a+r, \dots, a+(k-1)r$ are replaced by “polynomial progressions” $a + P_1(r), \dots, a + P_k(r)$. More precisely, let $\mathbb{Z}[\mathbf{m}]$ denote the ring of polynomials of one indeterminate variable \mathbf{m} with integer coefficients. Then Bergelson and Leibman [3] established the following polynomial version of Theorem 1:

Theorem 3 (Polynomial Szemerédi’s theorem). *Let $k \geq 1$, let $P_1, \dots, P_k \in \mathbb{Z}[\mathbf{m}]$ be such that $P_1(0) = \dots = P_k(0)$, let $\delta > 0$, and suppose that N is sufficiently large depending on $k, P_1, \dots, P_k, \delta$. Then any subset A of $[N]$ with cardinality $|A| \geq \delta N$ will contain at least one polynomial progression $a + P_1(r), a + P_2(r), \dots, a + P_k(r)$ with $r > 0$.*

Of course, Theorem 1 is the special case of Theorem 3 when $P_i = (i-1)\mathbf{m}$. As with Theorem 1, a partitioning argument shows that we may take $r \leq L(k, P_1, \dots, P_k, \delta)$ for some quantity L depending on the indicated parameters. The polynomial analogue of Theorem 2 was established by the authors in [19]:

Theorem 4 (Polynomial Szemerédi’s theorem in the primes). *Let $k \geq 2$, let $P_1, \dots, P_k \in \mathbb{Z}[\mathbf{m}]$ be such that $P_1(0) = \dots = P_k(0)$, $\varepsilon, \delta > 0$, and suppose that N is sufficiently large depending on $k, P_1, \dots, P_k, \delta, \varepsilon$. Then any subset A of $[N] \cap \mathbb{P}$ with $|A| \geq \delta |[N] \cap \mathbb{P}|$ will contain at least one polynomial progression $a + P_1(r), a + P_2(r), \dots, a + P_k(r)$ with $0 < r < N^\varepsilon$.*

In particular, this implies Theorem 2 with a bound $0 < r \leq N^{o(1)}$.

Remark 2. The condition $P_1(0) = \dots = P_k(0)$ in Theorem 3 can be relaxed to the property of *intersectivity* (that P_1, \dots, P_k have a common root in the profinite integers $\hat{\mathbb{Z}} = \lim_{\leftarrow m} \mathbb{Z}/m\mathbb{Z}$); see [4]. It is not yet known if Theorem 4 can similarly be relaxed to intersective polynomials, except in the $k = 2$ case which was established in [14]. We will not pursue this matter here.

1.2 Main new result

Our main result is to improve the bound on the r parameter in Theorem 4 to be polylogarithmic in size:

Theorem 5 (Short polynomial progressions in the primes). *Let $k \geq 2$, let $P_1, \dots, P_k \in \mathbb{Z}[\mathbf{m}]$ be such that $P_1(0) = \dots = P_k(0)$, $\varepsilon, \delta > 0$, and suppose that N is sufficiently large depending on $k, P_1, \dots, P_k, \delta, \varepsilon$. Then any subset A of $[N] \cap \mathbb{P}$ with $|A| \geq \delta |N| \cap \mathbb{P}|$ will contain at least one polynomial progression $a + P_1(r), a + P_2(r), \dots, a + P_k(r)$ with $0 < r < \log^L N$, where L depends only on $k, P_1, \dots, P_k, \delta$.*

In particular, there are infinitely polynomial progressions $a + P_1(r), a + P_2(r), \dots, a + P_k(r)$ consisting entirely of primes with $0 < r \ll \log^L a$, with L now depending only on k, P_1, \dots, P_k . This is new even in the case of arithmetic progressions $a, a + r, \dots, a + (k-1)r$.

A modification of the proof of Proposition 1 shows that some power of $\log N$ is needed in the upper bound on r in Theorem 5. However, we do not know what the optimal value of L is; our argument for general P_1, \dots, P_k uses the PET induction method [2], and as such L will grow rapidly with the degrees of the P_1, \dots, P_k . The dependence of L on δ occurs for technical reasons, and we conjecture that one can in fact select L to be independent of δ ; we can verify this conjecture in the arithmetic progression case $P_i = (i-1)\mathbf{m}$, and in fact we can take the explicit value $L := Ck2^k$ in this case for some fixed constant C (actually $C = 3$ would already suffice). We discuss this explicit variant of Theorem 5 in Section 7. Propositions 1, 2 suggest that we should in fact be able to set $L = k - 1 + \varepsilon$ in these cases, although our methods do not seem strong enough to achieve this, even in the $k = 3$ case.

It is possible that one might be able to directly modify of the arguments in [19] (which are in turn based on those in [10]) to establish Theorem 5; the main technical difficulties in doing so are the need of having to correlate the enveloping sieve with dual functions of unbounded functions and verifying the “correlation condition” required in that argument, in the setting when r is as small as $\log^L N$ (and the argument appears to have no chance of working when L is independent of δ). On the other hand, the need for the analogous correlation and boundedness conditions in [10] to prove Theorem 2 were recently removed² by Conlon, Fox, and Zhao [5], using a new method which they refer to as “densification”. We will be able to combine the densification method with the arguments in [19] to establish Theorem 5. As a consequence, we also obtain a slightly different proof of Theorem 4 than the one in [19], in which the (rather complicated) verification of the correlation condition is no longer necessary, but the densification arguments of Conlon, Fox, and Zhao are inserted instead.

² The “dual function condition” that the dual function of the enveloping sieve is bounded already failed in the arguments in [19], which was a significant cause of the complexity of that paper due to the need to find substitutes for this condition (in particular, the correlation condition became significantly more difficult to even state, let alone prove). But the arguments in [5] do not require any version of the dual function condition at all, leading to some simplifications in the current argument over those in [19].

2 Preliminary reductions

We now begin the proof of Theorem 5. We use the following asymptotic notation. We let N' be an asymptotic parameter tending to infinity along some sequence $N' = N'_j$ of natural numbers. All mathematical objects in this paper are implicitly permitted to depend on N' , unless they are explicitly designated to be *fixed*, in which case they are independent of N' . We use $X = O(Y)$, $X \ll Y$, or $Y \gg X$ to denote the estimate $|X| \leq CY$ for some fixed C , and $X = o(Y)$ to denote the estimate $|X| \leq c(N')Y$ where $c(N') \rightarrow 0$ as $N' \rightarrow \infty$. Our statements will be implicitly restricted to the regime in which N' is sufficiently large depending on all fixed parameters.

Suppose for sake of contradiction that Theorem 5 failed. Carefully negating the quantifiers (and relabeling N as N' , for reasons that will be clearer later), we conclude that we may find a fixed $k \geq 2$ and fixed polynomials $P_1, \dots, P_k \in \mathbb{Z}[\mathbf{m}]$ with $P_1(0) = \dots = P_k(0)$, fixed $\delta > 0$, a sequence $N' = N'_j$ of natural numbers going to infinity, and a set $A = A_{N'} \subset [N', 2N'] \cap \mathbb{P}$ such that

$$|A| \geq \delta |[N'] \cap \mathbb{P}|$$

and such that for any fixed $L > 0$, there are no polynomial progressions $a + P_1(r), a + P_2(r), \dots, a + P_k(r)$ in A with $0 < r < \log^L N'$ (recall that we assume N' sufficiently large depending on fixed quantities such as L).

The first few reductions are essentially the same to those in [19]. We begin with the “ W -trick” from [10] to eliminate local irregularities modulo small primes. We let $w = w_{N'}$ be a sufficiently slowly growing function of N' ; for instance, we could take $w := \frac{1}{10} \log \log \log N'$ as in [19] for sake of concreteness, although the precise value of w is unimportant. We then define the quantity

$$W := \prod_{p < w} p$$

and the natural number

$$N := \left\lfloor \frac{N'}{W} \right\rfloor.$$

From the prime number theorem³ we have

$$|A| \gg \frac{NW}{\log N}$$

and all elements of A larger than \sqrt{N} (say) are coprime to W . Thus, by the pigeon-hole principle, we may find $b \in [W]$ coprime to W (and depending on N' of course) such that

$$|\{n \in [N] : nW + b \in A\}| \gg \frac{NW}{\phi(W) \log N},$$

³ Actually, the weaker lower bound $\pi(x) \gg \frac{x}{\log x}$ of Chebyshev would suffice here.

where ϕ denotes the Euler totient function.

The domain $[N]$ is not quite translation invariant. In order to eliminate this (minor) difficulty, we will follow [19] and work instead in the cyclic group $X := \mathbb{Z}/N\mathbb{Z}$ using the obvious embedding $\iota : [N] \rightarrow X$. We give this space the uniform Haar probability measure, thus

$$\int_X f := \frac{1}{N} \sum_{x \in X} f(x)$$

for any $f : X \rightarrow \mathbb{R}$. We also define shift maps $T^h f : X \rightarrow \mathbb{R}$ for any $h \in \mathbb{Z}$ by

$$T^h f(x) := f(x + h);$$

clearly we have the identities

$$\begin{aligned} T^h T^k f &= T^{h+k} f \\ T^h(fg) &= (T^h f)(T^h g) \\ \int_X T^h f &= \int_X f \end{aligned}$$

for any $h, k \in \mathbb{Z}$ and $f, g : X \rightarrow \mathbb{R}$. We will use these identities frequently without further comment in the sequel.

We will need a fixed quantity $\varepsilon_0 > 0$ (depending only on k, P_1, \dots, P_k) to be chosen later. We define the function $f : X \rightarrow \mathbb{R}$ by the formula

$$f(\iota(n)) = \frac{\varepsilon_0}{10} \frac{\phi(W) \log N}{W} 1_{[N-\sqrt{N}] \setminus [\sqrt{N}]}(n) 1_A(nW + b) \quad (2)$$

for $n \in [N]$, where 1_A denotes the indicator function of A ; the reason for the normalizing factor $\frac{\varepsilon_0}{10}$ is so that f can be dominated by an enveloping sieve v , to be defined later. We then have

$$\int_X f \gg 1, \quad (3)$$

where we allow the implied constants here to depend on the fixed quantities δ and ε_0 .

Now let $L > 0$ be a sufficiently large fixed quantity (depending on $k, P_1, \dots, P_k, \delta$) to be chosen later. We will need the “coarse scale”

$$M := \log^L N \quad (4)$$

which will basically be the domain of interest for the polynomials P_1, \dots, P_k , and the “fine scale”

$$H := \log^{\sqrt{L}} N,$$

which will basically be the scale used for various applications of the van der Corput inequality. Note in particular that any quantity of size $O(H^{O(1)} \sqrt{M})$ will also be $o(M)$ if the implied constants do not depend on L , and L is large enough.

By hypothesis, if N (or N') is large enough, then the set A contains no polynomial progressions of the form $a + P_1(Wm), \dots, a + P_k(Wm)$ with $m \in [M]$. In particular, we see that

$$\Lambda(f, \dots, f) = 0 \quad (5)$$

where Λ is the k -linear form

$$\Lambda(f_1, \dots, f_k) := \mathbb{E}_{m \in [M]} \int_X T^{P_1(Wm)/W} f_1 \dots T^{P_k(Wm)/W} f_k \quad (6)$$

for k functions $f_1, \dots, f_k : X \rightarrow \mathbb{R}$, and we use the averaging notation

$$E_{a \in A} f(a) := \frac{1}{|A|} \sum_{a \in A} f(a).$$

Note that there are no “wraparound” issues caused by the embedding into $\mathbb{Z}/N\mathbb{Z}$, due to our removal in (2) of the elements $n \in [N]$ that are less than \sqrt{N} or larger than $N - \sqrt{N}$.

On the other hand, using the (multi-dimensional) polynomial Szemerédi theorem of Bergelson and Leibman [3], we have the following quantitative version of Theorem 3:

Theorem 6 (Quantitative Polynomial Szemerédi theorem). *Let $\delta > 0$ be fixed, and let $g : X \rightarrow \mathbb{R}$ obey the pointwise bounds*

$$0 \leq g \leq 1$$

(i.e., $0 \leq g(x) \leq 1$ for all $x \in X$), as well as the integral lower bound $\int_X g \geq \delta - o(1)$. Then we have

$$\Lambda(g, \dots, g) \geq c(\delta) - o(1)$$

where $c(\delta) > 0$ depends only on δ , k , and P_1, \dots, P_k .

Proof. See [19, Theorem 3.2]. In that theorem, the common value $P_1(0) = \dots = P_k(0)$ of the P_i was assumed to be zero, but the general case follows from this case by a simple change of variables.

This theorem cannot be directly applied to control $\Lambda(f, \dots, f)$, because f is not uniformly bounded. However, Theorem 5 will now be a consequence of the following claim.

Theorem 7 (Approximation by bounded function). *Suppose that $\varepsilon_0 > 0$ is sufficiently small (depending on k, P_1, \dots, P_k). Let $\varepsilon > 0$ be fixed. Suppose that L is a fixed quantity which is sufficiently large depending on $k, P_1, \dots, P_k, \varepsilon, \varepsilon_0$. Let f be as in (2). Then there exists $g : X \rightarrow \mathbb{R}$ with the pointwise bounds*

$$0 \leq g \leq 1,$$

such that

$$\left| \int_X f - \int_X g \right| \ll \varepsilon + o(1) \quad (7)$$

and

$$|\Lambda(f, \dots, f) - \Lambda(g, \dots, g)| \ll \varepsilon + o(1), \quad (8)$$

where the implied constants in the \ll notation do not depend on ε or L .

Let us assume Theorem 7 for now, and see how it implies Theorem 5. Let $\varepsilon_0 > 0$ be small enough for Theorem 7 to apply, $\varepsilon > 0$ be a sufficiently small fixed quantity (depending on δ, ε_0) to be chosen later, and let L be large enough depending on $k, P_1, \dots, P_k, \varepsilon, \varepsilon_0$ (in particular, L will depend on δ). Let g be as in Theorem 7, let M be defined by (4), and Λ defined by (6) (in particular, Λ depends on L). From (3), (7), and the triangle inequality we have

$$\int_X g \gg 1$$

if ε is small enough depending on ε_0, δ . Applying Theorem 6, we have

$$\Lambda(g, \dots, g) \gg c_0 - o(1)$$

for some $c_0 > 0$ depending on ε_0, δ but not on ε, L . By (8) and the triangle inequality, we thus have

$$\Lambda(f, \dots, f) > 0$$

if ε is small enough (depending on ε_0, δ) and N is large enough, contradicting (5); and Theorem 5 follows.

It remains to establish Theorem 7. This is the focus of the remaining sections of the paper.

Remark 3. The above arguments show that if one could remove the dependence of L on ε in Theorem 7, then one could also remove the dependence of L on δ in Theorem 5.

3 The enveloping sieve

As in [10], [19], we view f as a “positive density fraction” of a well-controlled *enveloping sieve* v , which is defined by the explicit formula

$$v(1(n)) := \frac{\phi(W) \log R}{W} \left(\sum_{d|Wn+b} \mu(d) \chi \left(\frac{\log d}{\log R} \right) \right)^2$$

for $n \in [N]$, where

$$R := N^{\varepsilon_0}, \quad (9)$$

μ is the Möbius function, and $\chi : \mathbb{R} \rightarrow \mathbb{R}$ is a fixed smooth even function supported on $[-1, 1]$ with the normalization

$$\int_0^1 |\chi'(t)|^2 dt = 1$$

(where χ' is the derivative of χ) and such that $\chi(0) \geq 1/2$ (say). By construction and (2), we have the pointwise bound

$$0 \leq f \leq v, \quad (10)$$

that is to say that $0 \leq f(n) \leq v(n)$ for all $n \in X$.

From [19, Corollary 10.5] we have the crude bound

$$\int_X \prod_{j=1}^J T^{h_j} v = 1 + o(1) + O \left(\text{Exp} \left(O \left(\sum_{1 \leq j < j' \leq J} \sum_{w \leq p \leq R^{\log R}, p | h_j - h_{j'}} \frac{1}{p} \right) \right) \right) \quad (11)$$

for any fixed J and any integers $h_1, \dots, h_J = O(\sqrt{N})$ (not necessarily distinct), assuming that ε_0 is sufficiently small depending on J , and where $\text{Exp}(x) := e^x - 1$, where the implied constant in the $O()$ exponent only depends on J . Here, of course, we use $p|n$ to denote the assertion that p divides n . In particular, we have the mean bound

$$\int_X v = 1 + o(1) \quad (12)$$

and the crude bound

$$\int_X T^{h_1} v \dots T^{h_J} v \ll \log^{O(1)} N \quad (13)$$

for any fixed J and any integers $h_1, \dots, h_J = O(\sqrt{N})$ (not necessarily distinct), assuming ε_0 is sufficiently small depending on J , and where the implied constant in the $O(1)$ exponent depends only on J .

One can use (11) to establish the following fundamental pseudorandomness property:

Proposition 3 (Polynomial forms condition). *Let $J, d, D \geq 1$ be fixed natural numbers. Suppose that ε_0 is sufficiently small depending on J, d, D , and that L is sufficiently large depending on J, d, D . Then for any polynomials $Q_1, \dots, Q_J \in \mathbb{Z}[\mathbf{m}_1, \dots, \mathbf{m}_d]$ of degree at most D , with coefficients of size $O(W^{O(1)})$, and with $Q_j - Q_{j'}$ non-constant for every $1 \leq j < j' \leq d$, and any convex body $\Omega \in [-M^2, M^2]^d$ of inradius⁴ at least $H^{1/2}$, we have the asymptotic*

$$\mathbb{E}_{\mathbf{h} \in \Omega \cap \mathbb{Z}^d} \int_X \prod_{j=1}^J T^{Q_j(\mathbf{h})} v = 1 + o(1). \quad (14)$$

⁴ The *inradius* of a convex body is the radius of the largest open ball one can inscribe inside the body.

This proposition was established in [19, Theorem 3.18] in the case when M is a small power of N ; the point is that M can be lowered to essentially $\log^L N$ for some large L . However, note that the number J of polynomials involved cannot be arbitrarily large depending on ε_0 . The main obstruction to reducing the size of the coarse scale M is that we need the “diagonal” contributions to (14) (such as those coming from the terms where one of the $Q_j(\mathbf{h})$ vanish) to be negligible when compared to the remaining terms. The sieve ν (or powers thereof, such as ν^2) tends to have size $\log^{O(1)} N$ on the average, and using this one expects to control diagonal contributions to (14) by something like $\log^{O(1)} N/M$, which will be negligible when M is a sufficiently large power of $\log N$.

Proof. We repeat the arguments from [19, §11]. For each \mathbf{h} , we see from (11) that

$$\int_X \prod_{j=1}^J T^{Q_j(\mathbf{h})} \nu = 1 + o(1) + O \left(\text{Exp} \left(O \left(\sum_{1 \leq j < j' \leq J} \sum_{w \leq p \leq R^{\log R}: p | Q_j(\mathbf{h}) - Q_{j'}(\mathbf{h})} \frac{1}{p} \right) \right) \right).$$

Thus it suffices to show that

$$\mathbb{E}_{\mathbf{h} \in \Omega \cap \mathbb{Z}^d} \text{Exp} \left(O \left(\sum_{1 \leq j < j' \leq J} \sum_{w \leq p \leq R^{\log R}: p | Q_j(\mathbf{h}) - Q_{j'}(\mathbf{h})} \frac{1}{p} \right) \right) = o(1).$$

Using the elementary bound $\text{Exp}(a+b) \ll \text{Exp}(2a) + \text{Exp}(2b)$ repeatedly, it suffices to show that

$$\mathbb{E}_{\mathbf{h} \in \Omega \cap \mathbb{Z}^d} \text{Exp} \left(O \left(\sum_{w \leq p \leq R^{\log R}: p | Q_j(\mathbf{h}) - Q_{j'}(\mathbf{h})} \frac{1}{p} \right) \right) = o(1).$$

for each $1 \leq j < j' \leq d'$.

We first dispose of the “globally bad” primes, in which p divides the entire polynomial $Q_j - Q_{j'}$. As $Q_j - Q_{j'}$ is non-constant and has coefficients $O(W^{O(1)})$, we see that the product of all such primes is $O(W^{O(1)})$. In [19, Lemma E.3], it is shown that

$$\sum_{p \geq w: p | n} \frac{1}{p} = o(1)$$

for any $n = O(W^{O(1)})$. Thus the contribution of such primes in the above sum is negligible.

In [19, Lemma E.1], it is shown that

$$\text{Exp} \left(O \left(\sum_{p \in A} \frac{1}{p} \right) \right) \ll \sum_{p \in A} \frac{\log^{O(1)} p}{p}$$

for any set A of primes. Thus it suffices to show that

$$\sum_{w \leq p \leq R^{\log R}} \frac{\log^{O(1)} p}{p} \mathbb{E}_{\mathbf{h} \in \Omega \cap \mathbb{Z}^d} 1_{p|Q_j(\mathbf{h}) - Q_{j'}(\mathbf{h}); p|Q_j - Q_{j'}} = o(1).$$

From [19, Lemma D.3], we conclude that if p does not divide $Q_j - Q_{j'}$, then the average $1_{p|Q_j(\mathbf{h}) - Q_{j'}(\mathbf{h})}$ on any cube in \mathbb{Z}^d of sidelength $1 \leq K \leq p$ is $O(\frac{1}{K})$. By [19, Corollary C.2], [19, Lemma C.4] and the inradius hypothesis we then have

$$\mathbb{E}_{\mathbf{h} \in \Omega \cap \mathbb{Z}^d} 1_{p|h_j - h_{j'}} \ll \frac{1}{p} + H^{-1/2}.$$

Thus we reduce to showing that

$$\sum_{w \leq p \leq R^{\log R}} \frac{\log^{O(1)} p}{p^2} + H^{-1/2} \frac{\log^{O(1)} p}{p} = o(1),$$

but this follows easily from Mertens' theorem (or the prime number theorem) and the definition of M and w , if L is large enough.

4 Averaged local Gowers norms

As in [19], we will control the left-hand side of (8) using some local Gowers norms, which we now define. Given any $d \geq 2$ and integers a_1, \dots, a_d and any scale $S \geq 1$, we define the *local Gowers uniformity norms* $U_S^{a_1, \dots, a_d}$ by the formula

$$\|f\|_{U_S^{a_1, \dots, a_d}}^{2^d} := \mathbb{E}_{m_1^{(0)}, \dots, m_d^{(0)}, m_1^{(1)}, \dots, m_d^{(1)} \in [S]} \int_X \prod_{(\omega_1, \dots, \omega_d) \in \{0,1\}^d} T^{m_1^{(\omega_1)} a_1 + \dots + m_d^{(\omega_d)} a_d} f \quad (15)$$

for $f : X \rightarrow \mathbb{R}$. Next, for any $t \geq 2$ and any d -tuple $\mathbf{Q} = (Q_1, \dots, Q_d)$ of polynomials $Q_i \in \mathbb{Z}[\mathbf{h}_1, \dots, \mathbf{h}_t, \mathbf{W}]$ in $t+1$ variables, we define the *averaged local Gowers uniformity norms* $U_S^{\mathbf{Q}([H]^t, W)}$ on functions $f : X \rightarrow \mathbb{R}$ by the formula

$$\|f\|_{U_S^{\mathbf{Q}([H]^t, W)}}^{2^d} := \mathbb{E}_{\mathbf{h} \in [H]^t} \|f\|_{U_S^{Q_1(\mathbf{h}, W), \dots, Q_d(\mathbf{h}, W)}}^{2^d}. \quad (16)$$

These are indeed norms; see [19, Appendix A]. One can extend these norms to complex-valued functions by inserting an alternating sequence of conjugation symbols in the product, but we will not need to use such an extension here. We remark that these expressions may also be defined for $d = 1$, but are merely seminorms instead of norms in that case.

From the Gowers-Cauchy-Schwarz inequality (see e.g. [11, Appendix B]) and Hölder's inequality, we record the useful inequality

$$\begin{aligned}
& \left| \mathbb{E}_{m_1^{(0)}, \dots, m_d^{(0)}, m_1^{(1)}, \dots, m_d^{(1)} \in [S]} \int_X \prod_{\omega \in \{0,1\}^d} T^{m_1^{(\omega_1)} a_1 + \dots + m_d^{(\omega_d)} a_d} f_\omega \right| \\
& \leq \prod_{\omega \in \{0,1\}^d} \|f_\omega\|_{U_S^{a_1, \dots, a_d}}^{2^d}
\end{aligned} \tag{17}$$

for any functions $f_\omega : X \rightarrow \mathbb{R}$ for $\omega \in \{0,1\}^d$ where we write $\omega := (\omega_1, \dots, \omega_d)$.

In a similar spirit, we have the following basic inequality:

Theorem 8 (Polynomial generalized von Neumann theorem). *Suppose that $\varepsilon_0 > 0$ is a fixed quantity which is sufficiently small depending on k, P_1, \dots, P_k , and that L is a fixed quantity which is sufficiently large depending on k, P_1, \dots, P_k . Then there exists fixed $t \geq 0, d \geq 2$ and a fixed d -tuple $\mathbf{Q} = (Q_1, \dots, Q_d)$ of polynomials $Q_i \in \mathbb{Z}[\mathbf{h}_1, \dots, \mathbf{h}_t, \mathbf{W}]$, none of which are identically zero, and which are independent of ε_0, L , such that one has the inequality*

$$|\Lambda(g_1, \dots, g_k)| \ll \min_{1 \leq i \leq k} \|g_i\|_{U_{\sqrt{M}}^{\mathbf{Q}([H]^t, W)}}^c + o(1)$$

for some fixed $c > 0$ and all $g_1, \dots, g_k : X \rightarrow \mathbb{R}$ obeying the pointwise bound $|g_i| \leq v + 1$ for all $i = 1, \dots, d$.

Proof. This is essentially [19, Theorem 4.5] (which was proven by a combination of PET induction, the Cauchy-Schwarz inequality, and the polynomial forms condition), with the only difference being that H and M are now polylogarithmic in N , rather than polynomial in N . However, an inspection of the proof of [19, Theorem 4.5] shows that this does not affect the arguments (after replacing the polynomial forms condition used there with Proposition 3, of course); the key relationships between H, M, N that are used in the proof are that $(HW)^{O(1)} \sqrt{M} = o(M)$ and that $(HWM)^{O(1)} = o(\sqrt{N})$, where the implied constants depend only on k, P_1, \dots, P_k (and in particular are independent of ε_0, L), and these properties are certainly obeyed for the choice of H and M used here. (The bound (13) is sufficient to deal with all the error terms arising from use of the van der Corput inequality in this regime.)

Setting $g_2 = \dots = g_k = 1$, we obtain in particular that

$$\left| \int_X g_1 \right| \ll \|g_1\|_{U_{\sqrt{M}}^{\mathbf{Q}([H]^t, W)}}^c + o(1).$$

(In fact, one can take $c = 1$ in this inequality by the standard monotonicity properties of the Gowers norms, see [19, Lemma A.3], but we will not need this improvement here.)

In view of this theorem, Theorem 7 is now a consequence of the following claim (after replacing ε with ε^c):

Theorem 9 (Approximation by bounded function, again). *Let $d \geq 2$ and $t \geq 0$ be fixed, and let $\mathbf{Q} = (Q_1, \dots, Q_d)$ be a fixed d -tuple of polynomials $Q_i \in$*

$\mathbb{Z}[\mathbf{h}_1, \dots, \mathbf{h}_t, \mathbf{W}]$, not identically zero. Let $\varepsilon_0 > 0$ be a fixed quantity that is sufficiently small depending on d, t, \mathbf{Q} . Let $\varepsilon > 0$ be fixed, and let L be a fixed quantity that is sufficiently large depending on $d, t, \mathbf{Q}, \varepsilon_0, \varepsilon$. Let \mathbf{v} be as above, and let $f : X \rightarrow \mathbb{R}$ obey the pointwise bound

$$0 \leq f \leq \mathbf{v}.$$

Then there exists $g : X \rightarrow \mathbb{R}$ with the pointwise bound

$$0 \leq g \leq 1,$$

such that

$$\|f - g\|_{U_{\sqrt{M}}^{\mathbf{Q}([H]^t, W)}} \ll \varepsilon + o(1). \quad (18)$$

It remains to establish Theorem 9. This is the objective of the remaining sections of the paper.

Remark 4. As before, if one could remove the dependence of L on ε in Theorem 9, then one could remove the dependence of L on δ in Theorem 5. Also, from this point on the number k of polynomials P_1, \dots, P_k in Theorem 5 no longer plays a role, and we will use the symbol k to denote other (unrelated) natural numbers.

5 The dense model theorem

Let d, t, \mathbf{Q} be as in Theorem 9. The averaged local Gowers norm $\|f\|_{U_{\sqrt{M}}^{\mathbf{Q}([H]^t, W)}}$ of a function $f : X \rightarrow \mathbb{R}$ can then be expressed by the identity

$$\|f\|_{U_{\sqrt{M}}^{\mathbf{Q}([H]^t, W)}}^{2^d} = \int f \mathcal{D}f \quad (19)$$

where the *dual function* $\mathcal{D}f = \mathcal{D}_{\sqrt{M}}^{\mathbf{Q}([H]^t, W)} f : X \rightarrow \mathbb{R}$ is defined by the formula

$$\begin{aligned} \mathcal{D}f := & \mathbb{E}_{\mathbf{h} \in [H]^t} \mathbb{E}_{m_1^{(0)}, \dots, m_d^{(0)}, m_1^{(1)}, \dots, m_d^{(1)} \in [\sqrt{M}]} \\ & \prod_{(\omega_1, \dots, \omega_d) \in \{0, 1\}^d \setminus \{0\}^d} T^{\sum_{i=1}^d (m_i^{(\omega_i)} - m_i^{(0)}) Q_i(\mathbf{h})} f \end{aligned} \quad (20)$$

More generally, we define

$$\begin{aligned} \mathcal{D}(f_\omega)_{\omega \in \{0, 1\}^d \setminus \{0\}^d} := & \mathbb{E}_{\mathbf{h} \in [H]^t} \mathbb{E}_{m_1^{(0)}, \dots, m_d^{(0)}, m_1^{(1)}, \dots, m_d^{(1)} \in [\sqrt{M}]} \\ & \prod_{(\omega_1, \dots, \omega_d) \in \{0, 1\}^d \setminus \{0\}^d} T^{\sum_{i=1}^d (m_i^{(\omega_i)} - m_i^{(0)}) Q_i(\mathbf{h})} f_\omega \end{aligned} \quad (21)$$

for any tuple of functions $f_\omega : X \rightarrow \mathbb{R}$ for $\omega \in \{0, 1\}^d \setminus \{0\}^d$.

Theorem 9 is then an immediate consequence of combining the following two theorems (with the function f appearing in Theorem 11 being replaced by the function $f - g$ appearing in Theorem 10).

Theorem 10 (Weak approximation by bounded function). *Let $d \geq 2$ and $t \geq 0$ be fixed, and let $\mathbf{Q} = (Q_1, \dots, Q_d)$ be a fixed d -tuple of polynomials $Q_i \in \mathbb{Z}[\mathbf{h}_1, \dots, \mathbf{h}_t, \mathbf{W}]$, not identically zero. Let $\varepsilon_0 > 0$ be a fixed quantity that is sufficiently small depending on d, t, \mathbf{Q} . Let $\varepsilon > 0$ be fixed, and let L be a fixed quantity that is sufficiently large depending on $d, t, \mathbf{Q}, \varepsilon_0, \varepsilon$. Let ν be as above, and let $f : X \rightarrow \mathbb{R}$ obey the pointwise bound*

$$0 \leq f \leq \nu.$$

Then there exists $g : X \rightarrow \mathbb{R}$ with the pointwise bound

$$0 \leq g \leq 1,$$

such that

$$\left| \int_X (f - g) \mathcal{D}(F_\omega)_{\omega \in \{0,1\}^d \setminus \{0\}^d} \right| \ll \varepsilon + o(1) \quad (22)$$

for all functions $F_\omega : X \rightarrow \mathbb{R}$ for $\omega \in \{0,1\}^d \setminus \{0\}^d$ with the pointwise bounds $-1 \leq F_\omega \leq 1$.

Theorem 11 (Densification). *Let $d \geq 2$ and $t \geq 0$ be fixed, and let $\mathbf{Q} = (Q_1, \dots, Q_d)$ be a fixed d -tuple of polynomials $Q_i \in \mathbb{Z}[\mathbf{h}_1, \dots, \mathbf{h}_t, \mathbf{W}]$, not identically zero. Let $\varepsilon_0 > 0$ be a fixed quantity that is sufficiently small depending on d, t, \mathbf{Q} . Let $\varepsilon > 0$ be fixed, and let L be a fixed quantity that is sufficiently large depending on $d, t, \mathbf{Q}, \varepsilon_0$. Let ν be as above, and let $f : X \rightarrow \mathbb{R}$ obey the pointwise bound*

$$|f| \leq \nu + 1,$$

and suppose that

$$\left| \int_X f \mathcal{D}(F_\omega)_{\omega \in \{0,1\}^d \setminus \{0\}^d} \right| \ll \varepsilon + o(1) \quad (23)$$

for all functions $F_\omega : X \rightarrow \mathbb{R}$ with the pointwise bounds $-1 \leq F_\omega \leq 1$. Then we have

$$\|f\|_{U_{\sqrt{M}}^{\mathbf{Q}([H]^t, \mathbf{W})}} \ll \varepsilon^c + o(1)$$

for some fixed $c > 0$ (independent of ε).

We prove Theorem 10 in this section, and Theorem 11 in the next section.

To prove Theorem 10, we invoke the dense model theorem, first established implicitly in [10] and then made more explicit in [19], [9], [16]. We use the formulation from [16, Theorem 1.1]:

Theorem 12 (Dense model theorem). *For every $\varepsilon > 0$, there is $K = (1/\varepsilon)^{O(1)}$ and $\varepsilon' > 0$ such that, whenever \mathcal{F} is a set of bounded functions from X to $[-1, 1]$, and $\nu : X \rightarrow \mathbb{R}^+$ obeys the bound*

$$\left| \int_X (\nu - 1) F_1 \dots F_{K'} \right| \leq \varepsilon'$$

for all $0 \leq K' \leq K$ and $F_1, \dots, F_{K'} \in \mathcal{F}$, and every function $f : X \rightarrow \mathbb{R}$ with $0 \leq f \leq \nu$, one has a function $g : X \rightarrow [0, 1]$ such that

$$\left| \int_X (f - g) F \right| \leq \varepsilon$$

for all $F \in \mathcal{F}$.

This reduces Theorem 10 to the following calculation:

Theorem 13 (Orthogonality to dual functions). *Let $d \geq 2$ and $t \geq 0$ be fixed, and let $\mathbf{Q} = (Q_1, \dots, Q_d)$ be a fixed d -tuple of polynomials $Q_i \in \mathbb{Z}[\mathbf{h}_1, \dots, \mathbf{h}_t, \mathbf{W}]$, not identically zero. Let $\varepsilon_0 > 0$ be a fixed quantity that is sufficiently small depending on d, t, \mathbf{Q} . Let $K > 0$ be a fixed integer, and let L be a fixed quantity that is sufficiently large depending on $d, t, \mathbf{Q}, \varepsilon_0, K$. Let ν be as above. Then one has*

$$\int_X (\nu - 1) \prod_{k=1}^K \mathcal{D}(F_{k,\omega})_{\omega \in \{0,1\}^d \setminus \{0\}^d} = o(1) \quad (24)$$

for all functions $F_{k,\omega} : X \rightarrow \mathbb{R}$ for $k = 1, \dots, K$, $\omega \in \{0,1\}^d \setminus \{0\}^d$ with the pointwise bounds $-1 \leq F_{k,\omega} \leq 1$.

Remark 5. The fact that L depends on K here is the sole reason why L depends on δ in Theorem 5 (note that no parameter related to δ or K appears in Theorem 11).

We now prove Theorem 13. Let $d, t, \mathbf{Q}, \varepsilon_0, K, L, \nu, F_{k,\omega}$ be as in that theorem. We expand out the left-hand side of (24) as the average of

$$\begin{aligned} & \int_X (\nu - 1) \mathbb{E}_{m_{i,k}^{(\omega)} \in [\sqrt{M}] \forall i=1, \dots, d; k=1, \dots, K; \omega=0,1} \\ & \prod_{(\omega_1, \dots, \omega_d) \in \{0,1\}^d \setminus \{0\}^d} \prod_{k=1}^K T^{\sum_{i=1}^d (m_{i,k}^{(\omega_i)} - m_i^{(0)}) Q_i(\mathbf{h}_k)} F_{k,\omega} \end{aligned} \quad (25)$$

as $\mathbf{h}_1, \dots, \mathbf{h}_K$ ranges over $[H]^t$.

We first deal with the degenerate cases in which $Q_i(\mathbf{h}_k) = 0$ for some i, k . By the Schwarz-Zippel lemma (see e.g. [19, Lemma D.3]), the number of tuples $(\mathbf{h}_1, \dots, \mathbf{h}_K)$ with this degeneracy is $O(H^{Kt-1})$. Meanwhile, from (12) and the boundedness of the $F_{k,\omega}$, each expression (25) is $O(1)$. Thus the total contribution of this case is $O(H^{-1})$, which is acceptable.

It thus suffices to show that the expression (25) is $o(1)$ uniformly for all $\mathbf{h}_1, \dots, \mathbf{h}_K$ with none of the $Q_i(\mathbf{h}_k)$ vanishing.

The next step is to “clear denominators” (as in [19]). Fix $\mathbf{h}_1, \dots, \mathbf{h}_K$, and write $D_i := \prod_{k=1}^K |Q_i(\mathbf{h}_k)|$ for $i = 1, \dots, d$. Then we have $1 \leq D_i \ll O(HW)^{O(K)}$, and we can write

$$D_i = Q_i(\mathbf{h}_k) r_{i,k}$$

for each $i = 1, \dots, d$, $k = 1, \dots, K$, and some $r_{i,k} = O(HW)^{O(K)}$.

Let $n_1^{(0)}, \dots, n_d^{(0)}, n_1^{(1)}, \dots, n_d^{(1)}$ be elements of $[M^{1/4}]$. Then if we shift each variable $m_{i,k}^{(\omega)}$ by $r_{i,k} n_i^{(\omega)}$, we can rewrite (25) as

$$\begin{aligned} & \int_X (\nu - 1) \mathbb{E}_{m_{i,k}^{(\omega)} \in [\sqrt{M}] - r_{i,k} n_i^{(\omega)} \forall i=1, \dots, d; k=1, \dots, K; \omega=0,1} \\ & \prod_{(\omega_1, \dots, \omega_d) \in \{0,1\}^d \setminus \{0\}^d} \prod_{k=1}^K T^{\sum_{i=1}^d (m_{i,k}^{(\omega_i)} - m_i^{(0)}) Q_i(\mathbf{h}_k) + (n_i^{(\omega_i)} - n_i^{(0)}) D_i} F_{k,\omega}. \end{aligned} \quad (26)$$

The shifted interval $[\sqrt{M}] - r_{i,k} n_i^{(\omega)}$ differs from $[\sqrt{M}]$ by shifts by a set of cardinality $O(M^{1/4}(HW)^{O(K)})$, and so by (13) one can replace the former by the latter after accepting an additive error of $O((\log^{O(1)} N) M^{1/4} (HW)^{O(1)} / \sqrt{M})$, where the implied constants in the exponents depend on K . It is at this point that we crucially use the hypothesis that L be large compared with K , to ensure that this error is still $o(1)$. Thus (26) can be written as

$$\begin{aligned} & \int_X (\nu - 1) \mathbb{E}_{m_{i,k}^{(\omega)} \in [\sqrt{M}] \forall i=1, \dots, d; k=1, \dots, K; \omega=0,1} \\ & \prod_{(\omega_1, \dots, \omega_d) \in \{0,1\}^d \setminus \{0\}^d} \prod_{k=1}^K T^{\sum_{i=1}^d (m_{i,k}^{(\omega_i)} - m_i^{(0)}) Q_i(\mathbf{h}_k) + (n_i^{(\omega_i)} - n_i^{(0)}) D_i} F_{k,\omega} \\ & + o(1). \end{aligned}$$

Averaging over all such $n_i^{(\omega)}$, we obtain

$$\begin{aligned} & \mathbb{E}_{m_{i,k}^{(\omega)} \in [\sqrt{M}] \forall i=1, \dots, d; k=1, \dots, K; \omega=0,1} \int_X (\nu - 1) \mathbb{E}_{n_1^{(0)}, \dots, n_d^{(0)}, n_1^{(1)}, \dots, n_d^{(1)} \in [M^{1/4}]} \\ & \prod_{(\omega_1, \dots, \omega_d) \in \{0,1\}^d \setminus \{0\}^d} \prod_{k=1}^K T^{\sum_{i=1}^d (m_{i,k}^{(\omega_i)} - m_i^{(0)}) Q_i(\mathbf{h}_k) + (n_i^{(\omega_i)} - n_i^{(0)}) D_i} F_{k,\omega} \\ & + o(1). \end{aligned}$$

Shifting the integral \int_X by $\sum_{i=1}^d (n_i^{(0)}) D_i$ and then using the Gowers-Cauchy-Schwarz inequality (17) (and the boundedness of the functions $\prod_{k=1}^K T^{\sum_{i=1}^d (m_{i,k}^{(\omega_i)} - m_i^{(0)}) Q_i(\mathbf{h}_k)} F_{k,\omega}$), we may bound this by

$$\|\nu - 1\|_{U_{M^{1/4}}^{D_1, \dots, D_d}} + o(1).$$

But from expanding out the Gowers norm (15) and using Proposition 3 to estimate the resulting 2^{2^d} terms (cf. [10, Lemma 5.2]), we see that

$$\|\nu - 1\|_{U_{M^{1/4}}^{D_1, \dots, D_d}}^{2^d} = o(1) \quad (27)$$

and Theorem 13 follows.

6 Densification

Now we prove Theorem 11. It will suffice to establish the following claim.

Proposition 4. *Let the notation and hypotheses be as in Theorem 11. Then one has*

$$\left| \mathbb{E}_{\mathbf{h} \in [H]^r} \mathbb{E}_{m_1^{(0)}, \dots, m_d^{(0)}, m_1^{(1)}, \dots, m_d^{(1)} \in [\sqrt{M}]} \int_X \prod_{\omega \in \{0,1\}^d} T^{m_1^{(\omega_1)} Q_1(\mathbf{h}) + \dots + m_d^{(\omega_d)} Q_d(\mathbf{h})} f_\omega \right| \ll \varepsilon^c + o(1) \quad (28)$$

for some fixed $c > 0$ (independent of ε), whenever $(f_\omega)_{\omega \in \{0,1\}^d}$ is a tuple of functions $f_\omega : X \rightarrow \mathbb{R}$ (with $\omega := (\omega_1, \dots, \omega_d)$), such that one of the f_ω is equal to f , and each of the remaining functions f_ω in the tuple either obey the pointwise bound $|f_\omega| \leq 1$ or $|f_\omega| \leq \nu$.

Indeed, given the above proposition, then by triangle inequality and decomposition we may replace the bounds $|f_\omega| \leq 1$ or $|f_\omega| \leq \nu$ with $|f_\omega| \leq \nu + 1$, and then by setting $f_\omega = f$ for every ω and using (19), we obtain the claim.

It remains to prove the proposition. We induct on the number of factors f_ω for which one has the bound $|f_\omega| \leq \nu$ instead of $|f_\omega| \leq 1$. First suppose that there are no such factors, thus $|f_\omega| \leq 1$ for all ω except for one ω , for which $f_\omega = f$. By permuting the cube $\{0,1\}^d$, we may assume that it is $f_{\{0\}^d}$ that is equal to f , with all other f_ω bounded in magnitude by 1. But then the expression in (28) may be rewritten as

$$\int_X f \mathcal{D}(f_\omega)_{\omega \in \{0,1\}^d \setminus \{0\}},$$

and the claim follows from the hypothesis (23).

Now suppose that at least one of the f_ω (other than the one equal to f) is bounded in magnitude by ν rather than 1. By permuting the cube we may assume that $|f_{\{0\}^d}| \leq \nu$. We then write the left-hand side of (28) as

$$\int_X f_{\{0\}^d} \mathcal{D}\mathbf{f}$$

where $\mathbf{f} := (f_\omega)_{\omega \in \{0,1\}^d \setminus \{0\}}$. By Cauchy-Schwarz, it thus suffices to show that

$$\int_X \nu (\mathcal{D}\mathbf{f})^2 \ll \varepsilon^c + o(1)$$

for some fixed $c > 0$. We will split this into two estimates,

$$|\int_X (\nu - 1)(\mathcal{D}\mathbf{f})^2| = o(1) \quad (29)$$

and

$$\int_X (\mathcal{D}\mathbf{f})^2 \ll \varepsilon^c + o(1). \quad (30)$$

We set aside (29) for now and work on (30). Bounding all the components of \mathbf{f} in magnitude by $v + 1$ and using Proposition 3, we see that

$$\int_X (\mathcal{D}\mathbf{f})^4 \ll 1,$$

so by Hölder's inequality, it suffices to show that

$$\int_X |\mathcal{D}\mathbf{f}| \ll \varepsilon^c + o(1)$$

(for a possibly different fixed $c > 0$). It thus suffices to show that

$$\left| \int_X g \mathcal{D}\mathbf{f} \right| \ll \varepsilon^c + o(1)$$

whenever $g : X \rightarrow \mathbb{R}$ is such that $|g| \leq 1$. But this expression is of the form (28) with $f_{\{0\}}$ replaced by g , and the claim then follows from the induction hypothesis.

It thus remains to show (29). We can rewrite

$$(\mathcal{D}\mathbf{f})^2 = \mathcal{D} \frac{\mathbf{Q} \oplus \mathbf{Q}([H]^t, W)}{\sqrt{M}} \mathbf{f}_2$$

where $\mathbf{Q} \oplus \mathbf{Q}$ is the $2d$ -tuple

$$\mathbf{Q} \oplus \mathbf{Q} = (Q_1, \dots, Q_d, Q_1, \dots, Q_d)$$

and $\mathbf{f}_2 = (f_{2,\omega})_{\omega \in \{0,1\}^{2d} \setminus \{0\}^d}$ is defined by setting

$$f_{2,\omega \oplus \{0\}^d} := f_\omega$$

and

$$f_{2,\{0\}^d \oplus \omega} := f_\omega$$

for $\omega \in \{0,1\}^d \setminus \{0\}^d$, and

$$f_{2,\omega \oplus \omega'} := 1$$

for $\omega, \omega' \in \{0,1\}^d \setminus \{0\}^d$. Applying the Gowers-Cauchy-Schwarz inequality (17), we may thus bound the left-hand side of (29) by

$$\|v - 1\|_{U \frac{\mathbf{Q} \oplus \mathbf{Q}([H]^t, W)}{\sqrt{M}}} \prod_{\omega \in \{0,1\}^d \setminus \{0\}^d} \|f_\omega\|_{U \frac{\mathbf{Q} \oplus \mathbf{Q}([H]^t, W)}{\sqrt{M}}}^2.$$

Bounding f_ω by v or 1 and using Proposition 3, we can bound

$$\|f_\omega\|_{U \frac{\mathbf{Q} \oplus \mathbf{Q}([H]^t, W)}{\sqrt{M}}} \ll 1$$

and further application of Proposition 3 (cf. (27)) gives

$$\|v - 1\|_{U_{\sqrt{M}}^{\mathbf{Q} \oplus \mathbf{Q}([H]^t, W)}} = o(1)$$

and the claim follows.

7 The linear case

We now explain why in the linear case $P_i = (i - 1)\mathbf{m}$ of Theorem 5, one may take L to be independent of δ . In the linear case, one can replace the averaged local Gowers norm $U_{\sqrt{M}}^{\mathbf{Q}([H]^t, W)}$ in Theorem 8 with the simpler norm $U_{\sqrt{M}}^{1, \dots, 1}$, where 1 appears $d = k - 1$ times; this follows by repeating the proof of [10, Proposition 5.3], after replacing some global averages with local ones. (In fact one could replace \sqrt{M} here by $M^{1-\sigma}$ for any fixed $\sigma > 0$.) As such, we can ignore the H parameter and the \mathbf{h} averaging, and just prove Theorem 13 in the case when $Q_1 = \dots = Q_d = 1$. Here, the “clearing denominators” step is unnecessary, and so L does not need to be large depending on K , which by Remark 5 ensures that the final L is independent of δ .

Remark 6. A more careful accounting of exponents (in particular, replacing (11) with a more precise asymptotic involving a singular series similar to that in (1)) allows one to take L as small as $Ck2^k$ for some absolute constant C ; we omit the details.

Acknowledgements The first author is supported by a Simons Investigator grant, the James and Carol Collins Chair, the Mathematical Analysis & Application Research Fund Endowment, and by NSF grant DMS-1266164. Part of this research was performed while the first author was visiting the Institute for Pure and Applied Mathematics (IPAM), which is supported by the National Science Foundation. The second author is supported by ISF grant 407/12. The second author was on sabbatical at Stanford while part of this work was carried out; she would like to thank the Stanford math department for its hospitality and support. Finally, the authors thank the anonymous referee for a careful reading of the paper.

References

1. J. Benatar, *The existence of small prime gaps in subsets of the integers*, preprint.
2. V. Bergelson, *Weakly mixing PET*, Ergodic Theory Dynam. Systems **7** (1987), no. 3, 337–349.
3. V. Bergelson, A. Leibman, *Polynomial extensions of van der Waerden’s and Szemerédi’s theorems*. J. Amer. Math. Soc. **9** (1996), no. 3, 725–753.
4. V. Bergelson, A. Leibman, E. Lesigne, *Intersective polynomials and the polynomial Szemerédi theorem*, Adv. Math. **219** (2008), no. 1, 369–388.
5. D. Conlon, J. Fox, Y. Zhao, *A relative Szemerédi theorem*, preprint.
6. D. Conlon, T. Gowers, *Combinatorial theorems in sparse random sets*, preprint.
7. P. X. Gallagher, *On the distribution of primes in short intervals*, Mathematika **23** (1976), 4–9.
8. D. Goldston, J. Pintz, C. Yıldırım, *Primes in tuples IV: Density of small gaps between consecutive primes*, Acta Arith. **160** (2013), no. 1, 37–53.

9. W. T. Gowers, *Decompositions, approximate structure, transference, and the Hahn-Banach theorem*, Bull. Lond. Math. Soc. **42** (2010), no. 4, 573–606.
10. B. Green, T. Tao, *The primes contain arbitrarily long arithmetic progressions*, Ann. of Math. (2) **167** (2008), no. 2, 481–547.
11. B. Green, T. Tao, *Linear equations in primes*, Annals of Math, **171** (2010), 1753–1850.
12. H. Halberstam, H.-E. Richert, *Sieve methods*, Academic Press, New York, 1974.
13. G. H. Hardy, J. E. Littlewood, *Some problems of “Partitio Numerorum”, III: On the expression of a number as a sum of primes*, Acta Math. **44** (1923), 1–70.
14. T. H. Le, *Intersective polynomials and the primes*, J. Number Theory **130** (2010), no. 8, 1705–1717.
15. J. Maynard, *Small gaps between primes*, preprint.
16. O. Reingold, L. Trevisan, M. Tulsiani and S. Vadhan, *Dense subsets of pseudorandom sets*, Electronic Colloquium on Computational Complexity, Proceedings of 49th IEEE FOCS, 2008.
17. M. Schacht, *Extremal results for random discrete structures*, preprint.
18. E. Szemerédi, *On sets of integers containing no k elements in arithmetic progression*, Acta Arith. **27** (1975), 299–345.
19. T. Tao, T. Ziegler, *The primes contain arbitrarily long polynomial progressions*, Acta Math. **201** (2008), no. 2, 213–305.
20. Y. Zhang, *Bounded gaps between primes*, to appear, Annals Math.
21. B. Zhou, *The Chen primes contain arbitrarily long arithmetic progressions*, Acta Arith. **138** (2009), no. 4, 301–315.